arXiv:1512.05808v1 [cs.LG] 17 Dec 2015

Successive Ray Refinement and Its Application to Coordinate Descent for LASSO

Jun Liu* SAS Institute Inc.
Zheng Zhao† Google Inc.
Ruiwen Zhang SAS Institute Inc.

December 21, 2015

Abstract

Coordinate descent is one of the most popular approaches for solving Lasso and its extensions due to its simplicity and efficiency. When applying coordinate descent to solving Lasso, we update one coordinate at a time while fixing the remaining coordinates. Such an update, which is usually easy to compute, greedily decreases the objective function value. In this paper, we aim to improve its computational efficiency by reducing the number of coordinate descent iterations. To this end, we propose a novel technique called Successive Ray Refinement (SRR). SRR makes use of the following ray continuation property on the successive iterations: for a particular coordinate, the value obtained in the next iteration almost always lies on a ray that starts at its previous iteration and passes through the current iteration. Motivated by this ray-continuation property, we propose that coordinate descent be performed not directly on the previous iteration but on a refined search point that has the following properties: on one hand, it lies on a ray that starts at a history solution and passes through the previous iteration, and on the other hand, it achieves the minimum objective function value among all the points on the ray. We propose two schemes for defining the search point and show that the refined search point can be efficiently obtained. Empirical results for real and synthetic data sets show that the proposed SRR can significantly reduce the number of coordinate descent iterations, especially for small Lasso regularization parameters.

1 Introduction

Lasso [12] is an effective technique for analyzing high-dimensional data. It has been applied successfully in various areas, such as machine learning, signal processing, image processing, medical imaging, and so on. Let $X = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p] \in \mathcal{R}^{n\times p}$ denote the data matrix composed of n samples with p variables, and let $\mathbf{y} \in \mathcal{R}^{n\times 1}$ be the response vector. In Lasso, we compute the $\boldsymbol{\beta}$ that optimizes

$$\min_{\boldsymbol{\beta}} f(\boldsymbol{\beta}) = \frac{1}{2}\|X\boldsymbol{\beta} - \mathbf{y}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_1, \tag{1}$$

*Corresponding author. E-mail address: junliu.nt@gmail.com.
†This work was done when Zheng Zhao was with SAS.

where the first term measures the discrepancy between the prediction and the response and the second term controls the sparsity of $\boldsymbol{\beta}$ with ℓ_1 regularization. The regularization parameter λ is nonnegative, and a larger λ usually leads to a sparser solution.

Researchers have developed many approaches for solving Lasso in Equation (1). Least Angle Regression (LARS) [3] is one of the most well-known homotopy approaches for Lasso. LARS adds or drops one variable at a time, generating a piecewise linear solution path for Lasso. Unlike LARS, other approaches usually solve Equation (1) according to some prespecified regularization parameters. These methods include the coordinate descent method [4, 18], the gradient descent method [1, 16], the interior-point method [6], the stochastic method [11], and so on. Among these approaches, coordinate descent is one of the most popular approaches due to its simplicity and efficiency. When applying coordinate descent to Lasso, we update one coordinate at a time while fixing the remaining coordinates. This type of update, which is easy to compute, can effectively decrease the objective function value in a greedy way.

To improve the efficiency of optimizing the Lasso problem in Equation (1), the screening technique has been extensively studied in [5, 8, 10, 13, 15, 19]. Screening 1) identifies and removes the variables that have zero entries in the solution $\boldsymbol{\beta}$ and 2) solves Equation (1) by using only the kept variables. When one is able to discard the variables that have zero entries in the final solution $\boldsymbol{\beta}$ and identify the signs of the nonzero entries, the Lasso problem in Equation (1) becomes a standard quadratic programming problem. However, it is usually very hard to identify all the zero entries, especially when the regularization parameter is small. In addition, the computational cost of Lasso usually increases as the the regularization parameter decreases. The computational cost increase motivates us to come up with an approach that can accelerate the computation of Lasso for small regularization parameters.

In this paper, we aim to improve the computational efficiency of coordinate descent by reducing its iterations. To this end, we propose a novel technique called Successive Ray Refinement (SRR). Our proposed SRR is motivated by an interesting ray-continuation property on the coordinate descent iterations: for a given coordinate, the value obtained in the next iteration almost always lies on a ray that starts at its previous iteration and passes through the current iteration. Figure 1 illustrates the ray-continuation property by using the data specified in Section 2. Motivated by this ray-continuation property, we propose that coordinate descent be performed not directly on the previous iteration but on a refined search point that has the following properties: on one hand, the search point lies on a ray that starts at a history solution and passes through the previous iteration, and on the other hand, the search point achieves the minimum objective function value among all the points on the ray. We propose two schemes for defining the search point, and we show that the refined search point can be efficiently computed. Experimental results on both synthetic and real data sets demonstrate that the proposed SRR can greatly accelerate the convergence of coordinate descent for Lasso, especially when the regularization parameter is small.

Organization The rest of this paper is organized as follows. We introduce the traditional coordinate descent for Lasso and present the ray-continuation property that motivates this paper in Section 2, propose the SRR technique in Section 3, discuss the efficient computation of the refinement factor that is used in SRR in Section 4, conduct

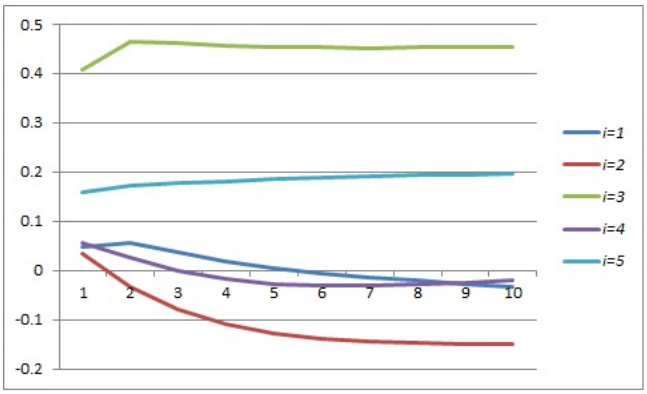

(a)

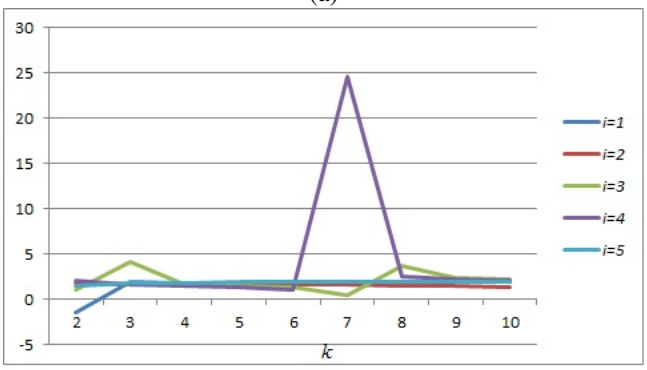

(b)

Figure 1: Illustration of the iterations of coordinate descent. For both plots, the x-axis corresponds to the iteration number k. The y-axis of plot (a) denotes β_i^k, the value of the ith coordinate in the kth iteration. The y-axis of plot (b) denotes α_i^k, which is computed using the equation $\beta_i^{k+1} = \alpha_i^k \beta_i^{k-1} + (1 - \alpha_i^k)\beta_i^k$. Ray-continuation property: for a given coordinate i, the value obtained in the next iteration denoted by β_i^{k+1} almost always lies on a ray that starts at its previous iteration, β_i^{k-1}, and passes through the current iteration, α_i^k. For numerical details of plot (a) and plot (b), see Table 1 and Table 2.

an eigenvalue analysis on the proposed SRR in Section 5, and compare SRR with related work in Section 6. We report experimental results on both synthetic and real data sets in Section 7, and we conclude this paper in Section 8.

Notations Throughout this paper, scalars are denoted by italic letters and vectors by bold face letters. Let $\|\cdot\|_1$ denote the ℓ_1 norm, let $\|\cdot\|_2$ denote the Euclidean norm, and let $\|\cdot\|_\infty$ denote the infinity norm. Let $\langle \mathbf{x}, \mathbf{y} \rangle$ denote the inner product between $\mathbf{x}$ and $\mathbf{y}$. Let a superscript denote the iteration number, and let a subscript denote the index of the variable or coordinate. We assume that X does not contain a zero column; that is, $\|\mathbf{x}_i\|_2 \neq 0, \forall i$.

2 Coordinate Descent For Lasso

In this section, we first review the coordinate descent method for solving Lasso, and then analyze the adjacent iterations to motivate the proposed SRR technique.

Let β_i^k denote the ith element of $\boldsymbol{\beta}$, which is obtained at the kth iteration of coordinate descent. In coordinate descent, we compute β_i^k while fixing $\beta_j = \beta_j^k, 1 \leq j < i$, and $\beta_j = \beta_j^{k-1}, i < j \leq p$. Specifically, β_i^k is computed as the minimizer to the following univariate optimization problem:

$$\beta_i^k = \arg\min_{\beta} f([\beta_1^k, \ldots, \beta_{i-1}^k, \beta, \beta_{i+1}^{k-1}, \ldots, \beta_p^{k-1}]^T).$$

It can be computed in a closed form as:

$$\beta_i^k = \frac{S(\mathbf{x}_i^T \mathbf{y} - \sum_{j<i} \mathbf{x}_i^T \mathbf{x}_j \beta_j^k - \sum_{j>i} \mathbf{x}_i^T \mathbf{x}_j \beta_j^{k-1}, \lambda)}{\|\mathbf{x}_i\|_2^2}, \tag{2}$$

where $S(\cdot, \cdot)$ is the shrinkage function

$$S(x, \lambda) = \begin{cases} x - \lambda & x > \lambda \\ x + \lambda & x < -\lambda \\ 0 & |x| \leq \lambda. \end{cases} \tag{3}$$

Let

$$\mathbf{r}_i^k = \mathbf{y} - X[\beta_1^k, \ldots, \beta_{i-1}^k, \beta_i^k, \beta_{i+1}^{k-1}, \ldots, \beta_p^{k-1}]^T \tag{4}$$

denote the residual obtained after updating β_i^{k-1} to β_i^k. With Equation (4), we can rewrite Equation (2) as

$$\beta_i^k = S(\beta_i^{k-1} + \frac{\mathbf{x}_i^T \mathbf{r}_{i-1}^k}{\|\mathbf{x}_i\|_2^2}, \frac{\lambda}{\|\mathbf{x}_i\|_2^2}). \tag{5}$$

In addition, with the updated β_i^k, we can update the residual from $\mathbf{r}_{i-1}^k$ to $\mathbf{r}_i^k$ as

$$\mathbf{r}_i^k = \mathbf{r}_{i-1}^k + \mathbf{x}_i(\beta_i^{k-1} - \beta_i^k). \tag{6}$$

Algorithm 1 illustrates solving Lasso via coordinate descent. Since the non-smooth ℓ_1 penalty in Equation (1) is separable, the algorithm is guaranteed to converge [14].

Algorithm 1 Coordinate Descent for Lasso

Input: X, $\mathbf{y}$, λ
Output: $\boldsymbol{\beta}^k$

$k = 0$, $\boldsymbol{\beta}^0 = \mathbf{0}$, $\mathbf{r}^0 = \mathbf{y}$
repeat
  Set $k = k + 1$, $\mathbf{r}^k = \mathbf{r}^{k-1}$
  for $i = 1$ to p **do**
    Compute $\beta_i^k = S(\beta_i^{k-1} + \frac{\mathbf{x}_i^T \mathbf{r}^k}{\|\mathbf{x}_i\|_2^2}, \frac{\lambda}{\|\mathbf{x}_i\|_2^2})$
    Update residual $\mathbf{r}^k = \mathbf{r}^k + \mathbf{x}_i(\beta_i^{k-1} - \beta_i^k)$
  end for
until convergence criterion satisfied

Table 1: Applying coordinate descent in Algorithm 1 to solving Lasso in Equation (1) with $\lambda = 0$. Since $\lambda = 0$ and X is invertible, the optimal function value is 0.

k	β_1^k	β_2^k	β_3^k	β_4^k	β_5^k	$f(\boldsymbol{\beta}^k)$
1	0.048912	0.034041	0.407960	0.055687	0.160413	0.052449
2	0.057182	-0.033692	0.465254	0.027810	0.171740	0.017591
3	0.036909	-0.079955	0.463604	-0.000612	0.177708	0.008085
4	0.019050	-0.108954	0.458440	-0.017618	0.182115	0.003933
5	0.005418	-0.126712	0.455218	-0.026135	0.185698	0.002304
6	-0.005122	-0.137403	0.453694	-0.029295	0.188740	0.001653
7	-0.013567	-0.143688	0.453262	-0.029210	0.191398	0.001358
8	-0.020585	-0.147239	0.453491	-0.027216	0.193771	0.001187
9	-0.026604	-0.149101	0.454104	-0.024144	0.195923	0.001060
10	-0.031899	-0.149927	0.454929	-0.020507	0.197895	0.000950
...						
28	-0.081090	-0.142355	0.467964	0.030835	0.217715	0.000106
29	-0.082490	-0.142044	0.468368	0.032406	0.218288	0.000093
30	-0.083806	-0.141752	0.468749	0.033883	0.218827	0.000082
...						
100	-0.103999	-0.137267	0.474584	0.056543	0.227099	1.3349e-08
101	-0.104015	-0.137264	0.474589	0.056561	0.227105	1.1785e-08
102	-0.104030	-0.137261	0.474593	0.056577	0.227111	1.0403e-08
103	-0.104044	-0.137258	0.474597	0.056593	0.227117	9.1839e-09
104	-0.104057	-0.137255	0.474601	0.056608	0.227122	8.1074e-09
105	-0.104069	-0.137252	0.474604	0.056621	0.227127	7.1571e-09

We demonstrate Algorithm 1 using the following randomly generated X and $\mathbf{y}$:

$$X = \begin{bmatrix} -0.204708 & 0.478943 & -0.519439 & -0.555730 & 1.965781 \\ 1.393406 & 0.092908 & 0.281746 & 0.769023 & 1.246435 \\ 1.007189 & -1.296221 & 0.274992 & 0.228913 & 1.352917 \\ 0.886429 & -2.001637 & -0.371843 & 1.669025 & -0.438570 \\ -0.539741 & 0.476985 & 3.248944 & -1.021228 & -0.577087 \end{bmatrix}, \quad (7)$$

$$\mathbf{y} = [0.124121, 0.302614, 0.523772, 0.000940, 1.343810]^T. \quad (8)$$

We show the iterations of coordinate descent for Lasso with $\lambda = 0$ in Table 1 and Figure 1 (a). We set $\lambda = 0$ to facilitate the eigenvalue analysis in Section 5. Note that the results reported here also generalize to Lasso, because if we know the sign of the optimal solution $\boldsymbol{\beta}^*$, the nonzero entries of $\boldsymbol{\beta}^*$ can be solved by the following equivalent convex smooth problem:

$$\min_{\boldsymbol{\beta}} \frac{1}{2} \| \sum_{i:s_i \neq 0} \mathbf{x}_i \beta_i - \mathbf{y} \|_2^2 + \lambda \sum_{i:s_i \neq 0} \beta_i s_i, \quad (9)$$

where $s_i = 0$ if $\beta_i^* = 0$, $s_i = 1$ if $\beta_i^* > 0$, and $s_i = -1$ if $\beta_i^* < 0$.

It can be observed from the results in Table 1 and Figure 1 (a) that we can obtain an approximate solution with a small objective function value within a few iterations. However, achieving a solution with high precision takes quite a few iterations for this example. More interestingly, for a particular coordinate, the value obtained in the next iteration almost always lies on a ray that starts at its previous iteration and passes through the current iteration. To show this, we compute α_i^k that satisfies the following equation:

$$\beta_i^{k+1} = \alpha_i^k \beta_i^{k-1} + (1 - \alpha_i^k) \beta_i^k. \quad (10)$$

Table 2 and Figure 1 (b) show the values of α_i^k for different iterations. It can be observed that the values of α_i^k are almost always positive except α_1^2 for this example. In addition, most of the values of α_i^k are larger than 1. We tried quite a few synthetic data and observed a similar phenomenon.

For a particular iteration number k, if $\alpha_i^k = \alpha, \forall i$, we can easily achieve $\boldsymbol{\beta}^{k+1} = \alpha \boldsymbol{\beta}^{k-1} + (1 - \alpha) \boldsymbol{\beta}^k$ without needing to perform any coordinate descent iteration. This motivated us to come up with the successive ray refinement technique to be discussed in the next section.

3 Successive Ray Refinement

In the proposed SRR technique, we make use of the ray-continuation property shown in Figure 1, Table 1, and Table 2. Our idea is as follows: To obtain $\boldsymbol{\beta}^{k+1}$, we perform coordinate descent based on a refined search point $\mathbf{s}^k$ rather than on its previous solution $\boldsymbol{\beta}^k$. We propose setting the refined search point as:

$$\mathbf{s}^k = (1 - \alpha^k) \mathbf{h}^k + \alpha^k \boldsymbol{\beta}^k, \quad (11)$$

Table 2: Illustration of the ray-continuation property $\beta_i^{k+1} = \alpha_i^k \beta_i^{k-1} + (1 - \alpha_i^k)\beta_i^k$ based on the results obtained in Table 1. All the α_i^k are positive except α_1^2.

k	α_1^k	α_2^k	α_3^k	α_4^k	α_5^k
2	-1.451503	1.683019	0.971191	2.019562	1.526899
3	1.880924	1.626844	4.128185	1.598324	1.738268
4	1.763341	1.612353	1.624041	1.500861	1.813212
5	1.773143	1.602012	1.473119	1.370962	1.849048
6	1.801288	1.587906	1.283008	0.973180	1.873809
7	1.830983	1.564952	0.469896	24.529117	1.892594
8	1.857679	1.524402	3.679312	2.540794	1.906640
9	1.879781	1.443912	2.346272	2.183873	1.916832
10	1.897021	1.240529	2.128716	2.068764	1.924043
...					
28	1.939534	1.939949	1.939646	1.939628	1.939556
29	1.939545	1.939821	1.939620	1.939608	1.939560
30	1.939552	1.939736	1.939602	1.939594	1.939562

where $\mathbf{h}^k$ is a properly chosen history solution, $\boldsymbol{\beta}^k$ is the current solution, and α^k is an optimal refinement factor that optimizes the following univariate optimization problem:

$$\alpha^k = \arg\min_{\alpha}\{g(\alpha) = f((1-\alpha)\mathbf{h}^k + \alpha\boldsymbol{\beta}^k)\}. \tag{12}$$

The setting of $\mathbf{h}^k$ to one of the history solutions is based on the following two considerations. First, we aim to use the ray-continuation property to reduce the number of iterations. Second, we need to ensure that the univariate optimization problem in Equation (12) can be efficiently computed. We discuss the computation of Equation (12) in Section 4.

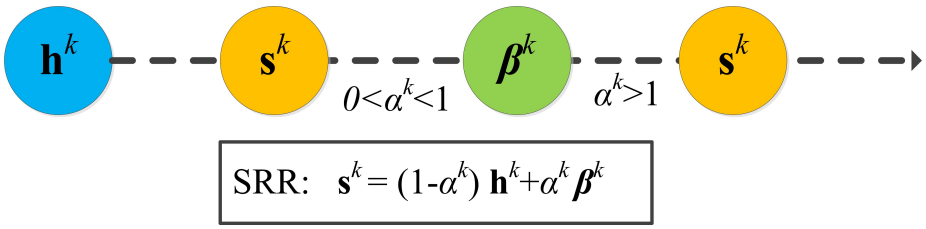

Figure 2: The proposed SRR technique. The search point $\mathbf{s}^k$ lies on the ray that starts from a properly chosen history solution $\mathbf{h}^k$ and passes through the current solution $\boldsymbol{\beta}^k$, and meanwhile it achieves the minimum objective function value among all the points on the ray, optimizing Equation (12).

Figure 2 illustrates the proposed SRR technique. When $\alpha^k = 1$, we have $\mathbf{s}^k = \boldsymbol{\beta}^k$; that is, the refined search point becomes the current solution $\boldsymbol{\beta}^k$. When $\alpha^k = 0$, we

have $\mathbf{s}^k = \mathbf{h}^k$; that is, the refined search point becomes the specified history solution $\mathbf{h}^k$. However, our next theorem shows that $\mathbf{s}^k \neq \mathbf{h}^k$ because α^k is always positive. In other words, the search point always lies on a ray that starts with the history point $\mathbf{h}^k$ and passes through the current solution $\boldsymbol{\beta}^k$.

Theorem 1 *Assume that the history point* $\mathbf{h}^k$ *satisfies*

$$f(\mathbf{h}^k) > f(\boldsymbol{\beta}^k). \tag{13}$$

Then, α^k *that minimizes Equation* (12) *is positive. In addition, if* $X\mathbf{h}^k \neq X\boldsymbol{\beta}^k$, α^k *is unique.*

Proof It is easy to verify that $g(\alpha)$ is convex. Therefore, α^k that minimizes Equation (12) has at least one solution. Equation (13) leads to

$$g(1) < g(0). \tag{14}$$

Therefore, the global refinement factor $\alpha^k \neq 0$. Next, we show that α^k cannot be negative.

If $\alpha^k < 0$, due to the convexity of $g(\alpha)$, we have

$$g((1-\theta)\alpha^k + \theta) \leq (1-\theta)g(\alpha^k) + \theta g(1), \forall \theta \in [0,1]. \tag{15}$$

Setting $\theta = \frac{\alpha^k}{\alpha^k - 1}$, we have

$$g(0) \leq \frac{-1}{\alpha^k - 1} g(\alpha^k) + \frac{\alpha^k}{\alpha^k - 1} g(1), \forall \theta \in [0,1]. \tag{16}$$

Making use of Equation (14), we have $g(1) < g(\alpha^k)$. This contradicts the fact that α^k minimizes Equation (12). Therefore, α^k is always positive.

If $X\mathbf{h}^k \neq X\boldsymbol{\beta}^k$, $g(\alpha)$ is strongly convex and thus α^k is unique. This ends the proof of this theorem. □

For coordinate descent, the condition in Equation (13) always holds, because the objective function value keeps decreasing. The selection of an appropriate $\mathbf{h}^k$ is key to the success of the proposed SRR, and the following theorem says that if $\mathbf{h}^k$ is good enough, the refined search solution $\mathbf{s}^k$ is an optimal solution to Equation (1).

Theorem 2 *Let* $\boldsymbol{\beta}^*$ *be an optimal solution to Equation* (1)*. If*

$$\boldsymbol{\beta}^* - \mathbf{h}^k = \gamma(\boldsymbol{\beta}^k - \mathbf{h}^k), \tag{17}$$

for some positive γ, $\mathbf{s}^k$ *achieved by SRR in Equation* (11) *satisfies* $f(\mathbf{s}^k) = f(\boldsymbol{\beta}^*)$.

Proof When setting $\alpha^k = \gamma$, we have $\mathbf{s}^k = \boldsymbol{\beta}^*$ under the assumption in Equation (17). Therefore, with the SRR technique, we can obtain a refined solution $\mathbf{s}^k$ that is an optimal solution to Equation (1). □

In the following subsections, we discuss two schemes for choosing the history solution $\mathbf{h}^k$.

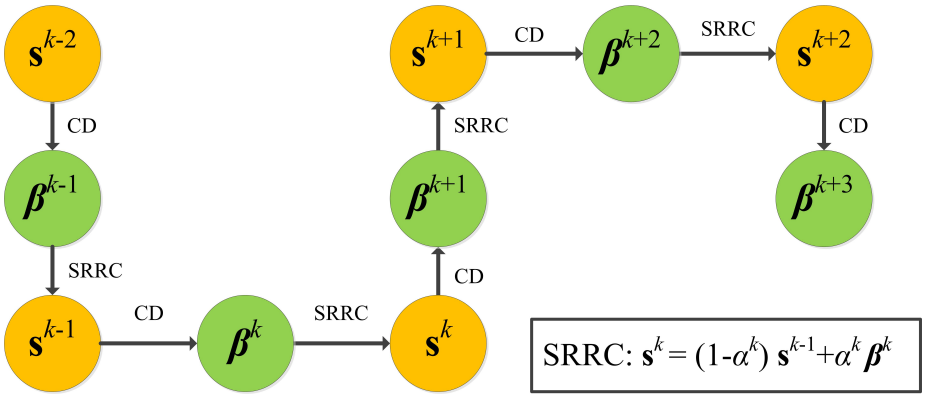

Figure 3: Illustration of the proposed SRRC technique. $\mathbf{s}^k$ is the refined search point and $\boldsymbol{\beta}^{k+1}$ is the point that is obtained by applying coordinate descent (CD) based on the refined search point $\mathbf{s}^k$. In this illustration, it is assumed that the optimal refinement factor α^k is larger than 1. When $\alpha^k \in (0, 1)$, $\mathbf{s}^k$ lies between $\mathbf{s}^{k-1}$ and $\boldsymbol{\beta}^k$.

3.1 Successive Ray Refinement Chain

In the first scheme, we set

$$\mathbf{h}^k = \mathbf{s}^{k-1}. \tag{18}$$

That is, the history point is set to the most recent refined search point. Figure 3 demonstrates this scheme. Since the generated points follow a chain structure, we call this scheme the Successive Ray Refinement Chain (SRRC). In SRRC, $\mathbf{s}^{k-1}$, $\boldsymbol{\beta}^k$, and $\mathbf{s}^k$ lie on the same line. In addition, coordinate descent (CD) controls the direction of the chain. In this illustration, it is assumed that the optimal refinement factor α^k is larger than 1 in each step. According to Theorem 1, $\alpha^k > 0$. When $\alpha^k \in (0, 1)$, $\mathbf{s}^k$ lies between $\mathbf{s}^{k-1}$ and $\boldsymbol{\beta}^k$. When $\alpha^k = 1$, $\mathbf{s}^k$ coincides with $\boldsymbol{\beta}^k$.

In Algorithm 2, we apply the proposed SRRC to coordinate descent for Lasso. Compared with the traditional coordinate descent in Algorithm 1, the coordinate update is based on the search point $\mathbf{s}^{k-1}$ rather than on the previous solution $\boldsymbol{\beta}^{k-1}$. When α^k in line 9 of Algorithm 2 is set to 1, Algorithm 2 becomes identical to Algorithm 1.

Table 3 illustrates Algorithm 2 with the same input X and $\mathbf{y}$ that are used in Table 1. Comparing Table 3 with Table 1, we can see that the number of iterations can be significantly reduced with the usage of the SRRC technique. Specifically, to achieve a function value of below 10^{-3}, the traditional coordinate descent takes 10 iterations, whereas the one with the SRRC technique takes 7 iterations; to achieve a function value below 10^{-4}, the traditional coordinate descent takes 29 iterations, whereas the one with the SRRC technique 14 iterations; and to achieve a function value below 10^{-8}, the traditional coordinate descent takes 103 iterations, whereas the one with the SRRC technique takes 16 iterations.

As can be seen from Figure 3, we generate two sequences: $\{\mathbf{s}^k\}$ and $\{\boldsymbol{\beta}^k\}$. At iteration k, the SRRC technique is very greedy in that it constructs the search point $\mathbf{s}^k$ by using the two existing points $\mathbf{s}^{k-1}$ and $\boldsymbol{\beta}^k$ to achieve the lowest objective function

Algorithm 2 Coordinate Descent plus SRRC (CD+SRRC) for Lasso

Input: X, $\mathbf{y}$, λ
Output: $\boldsymbol{\beta}^k$

Set $k = 0$, $\mathbf{s}^0 = \mathbf{0}$, $\mathbf{r}_s^0 = \mathbf{y}$
repeat
  Set $k = k + 1$, $\mathbf{r}^k = \mathbf{r}_s^{k-1}$
  for $i = 1$ to p **do**
    Compute $\beta_i^k = S(s_i^{k-1} + \frac{\mathbf{x}_i^T \mathbf{r}^k}{\|\mathbf{x}_i\|_2^2}, \frac{\lambda}{\|\mathbf{x}_i\|_2^2})$
    Obtain $\mathbf{r}^k = \mathbf{r}^k + \mathbf{x}_i(s_i^{k-1} - \beta_i^k)$
  end for
  if convergence criterion not satisfied **then**
    Set $\alpha^k = \arg\min_\alpha f((1-\alpha)\mathbf{s}^{k-1} + \alpha\boldsymbol{\beta}^k)$
    Set $\mathbf{s}^k = (1-\alpha^k)\mathbf{s}^{k-1} + \alpha^k\boldsymbol{\beta}^k$
    Set $\mathbf{r}_s^k = (1-\alpha^k)\mathbf{r}_s^{k-1} + \alpha^k\mathbf{r}^k$
  end if
until convergence criterion satisfied

Table 3: Illustration of coordinate descent plus SRRC for solving the same problem as in Table 1. Note that the optimal function value is 0.

k	β_1^k	β_2^k	β_3^k	β_4^k	β_5^k	$f(\boldsymbol{\beta}^k)$	α^k
1	0.048912	0.034041	0.407960	0.055687	0.160413	0.052449	
2	0.058130	-0.041464	0.471828	0.024612	0.173040	0.016773	1.114740
3	0.022324	-0.108065	0.459034	-0.018702	0.181180	0.004209	1.520601
4	-0.000996	-0.137517	0.452455	-0.033262	0.187045	0.001791	1.610933
5	-0.010602	-0.144776	0.452250	-0.033032	0.190141	0.001452	1.114831
6	-0.029911	-0.153851	0.453133	-0.026748	0.196740	0.001091	2.700667
7	-0.047531	-0.149751	0.458275	-0.006750	0.203971	0.000632	3.936469
8	-0.052347	-0.148709	0.459668	-0.001392	0.205944	0.000530	1.398237
9	-0.058803	-0.147299	0.461526	0.005839	0.208586	0.000407	2.059226
10	-0.064683	-0.145997	0.463220	0.012415	0.210994	0.000308	2.134921
...							
13	-0.078615	-0.142905	0.467249	0.028058	0.216701	0.000129	1.100414
14	-0.093529	-0.139593	0.471552	0.044782	0.222808	0.000023	9.617055
15	-0.094167	-0.139451	0.471743	0.045510	0.223071	0.000020	0.997764
16	-0.104249	-0.137213	0.474657	0.056823	0.227201	3.3020e-11	16.530123
...							
28	-0.104260	-0.137210	0.474660	0.056835	0.227205	6.3207e-15	1.197748
29	-0.104260	-0.137210	0.474660	0.056835	0.227205	5.7081e-15	0.890278
30	-0.104260	-0.137210	0.474660	0.056835	0.227205	2.0807e-15	6.945350

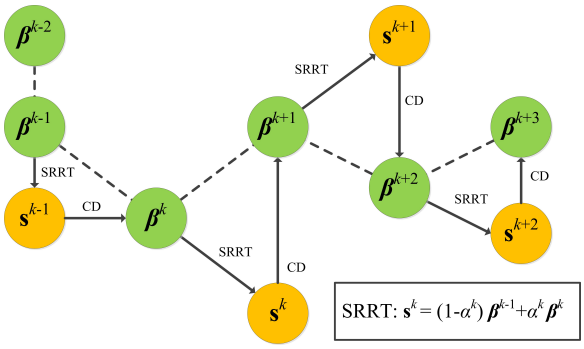

Figure 4: Illustration of the proposed SRRT technique. $\mathbf{s}^k$ is the refined search point and $\boldsymbol{\beta}^{k+1}$ is the point obtained by applying coordinate descent (CD) based on the refined search point $\mathbf{s}^k$. In this illustration, it is assumed that the optimal refinement factor α^k is larger than 1. When $\alpha^k \in (0, 1)$, $\mathbf{s}^k$ lies between $\boldsymbol{\beta}^{k-1}$ and $\boldsymbol{\beta}^k$.

value. If the search point $\mathbf{s}^{k-1}$ is dense at some iteration number k and $\alpha^k \neq 1$, it can be shown that $\mathbf{s}^k$ is also dense. This is not good for Lasso, which usually has a sparse solution. Interestingly, our empirical simulations show that Algorithm 2 can set $\alpha^k = 1$ in some iterations, leading to a sparse search point.

3.2 Successive Ray Refinement Triangle

In the second scheme, we set

$$\mathbf{h}^k = \boldsymbol{\beta}^{k-1}. \tag{19}$$

Figure 4 demonstrates this scheme. Since the generated points follow a triangle structure, we call this scheme the Successive Ray Refinement Triangle (SRRT). SRRT is less greedy compared to SRRC because $\boldsymbol{\beta}^{k-1}$ leads to a higher objective function value than $\mathbf{s}^{k-1}$ leads to. However, SRRT can sometimes outperform SRRC in solving Lasso.

Algorithm 3 shows the application of the proposed SRRT technique to coordinate descent for Lasso. Similar to Algorithm 2, if α^k in line 9 is set to 1, Algorithm 3 reduces to the traditional coordinate descent in Algorithm 1. Table 4 illustrates Algorithm 3. Similar to SRRC, SRRT greatly reduces the number of iterations used in coordinate descent for Lasso.

3.3 Convergence of CD plus SRR

In this subsection, we show that both the combination of CD and SRRC (CD+SRRC) and the combination of CD and SRRT (CD+SRRT) are guaranteed to converge.

Theorem 3 *For the sequence* $\mathbf{s}^0, \boldsymbol{\beta}^1, \mathbf{s}^1, \boldsymbol{\beta}^2, \mathbf{s}^2, \boldsymbol{\beta}^3, \ldots$ *generated by CD+SRRC and CD+SRRT, the objective function value is monotonically decreasing until convergence; that is,*

$$f(\mathbf{s}^{k-1}) \geq f(\boldsymbol{\beta}^k) \geq f(\mathbf{s}^k) \geq f(\boldsymbol{\beta}^{k+1}). \tag{20}$$

Algorithm 3 Coordinate Descent plus SRRT (CD+SRRT) for Lasso

Input: X, $\mathbf{y}$, λ
Output: $\boldsymbol{\beta}^k$

Set $k = 0$, $\mathbf{s}^0 = \mathbf{0}$, $\mathbf{r}_s^0 = \mathbf{y}$
repeat
  Set $k = k + 1$, $\mathbf{r}^k = \mathbf{r}_s^{k-1}$
  for $i = 1$ to p **do**
    Compute $\beta_i^k = S(s_i^{k-1} + \frac{\mathbf{x}_i^T \mathbf{r}^k}{\|\mathbf{x}_i\|_2^2}, \frac{\lambda}{\|\mathbf{x}_i\|_2^2})$
    Obtain $\mathbf{r}^k = \mathbf{r}^k + \mathbf{x}_i(s_i^{k-1} - \beta_i^k)$
  end for
  if convergence criterion not satisfied **then**
    Set $\alpha^k = \arg\min_\alpha f((1-\alpha)\boldsymbol{\beta}^{k-1} + \alpha\boldsymbol{\beta}^k)$
    Set $\mathbf{s}^k = (1-\alpha^k)\boldsymbol{\beta}^{k-1} + \alpha^k\boldsymbol{\beta}^k$
    Set $\mathbf{r}_s^k = (1-\alpha^k)\mathbf{r}^{k-1} + \alpha^k\mathbf{r}^k$
  end if
until convergence criterion satisfied

Table 4: Illustration of coordinate descent plus SRRT for solving the same problem as in Table 1. Note that the optimal function value is 0.

k	β_1^k	β_2^k	β_3^k	β_4^k	β_5^k	$f(\boldsymbol{\beta}^k)$	α^k
1	0.048912	0.034041	0.407960	0.055687	0.160413	0.052449	
2	0.058130	-0.041464	0.471828	0.024612	0.173040	0.016773	1.114740
3	0.032838	-0.089244	0.463272	-0.006319	0.178907	0.006746	1.077199
4	-0.010078	-0.154209	0.449373	-0.043957	0.189153	0.001610	2.336008
5	-0.015176	-0.152741	0.450482	-0.038189	0.191151	0.001435	0.960038
6	-0.087427	-0.134480	0.471220	0.044793	0.220728	0.000061	15.373834
7	-0.098214	-0.134199	0.474324	0.054977	0.225125	0.000019	1.138984
8	-0.104044	-0.135210	0.475348	0.058970	0.227325	0.000005	1.492143
9	-0.106491	-0.136042	0.475553	0.060121	0.228188	0.000002	1.414778
10	-0.106739	-0.136361	0.475482	0.059962	0.228250	0.000002	1.141313
...							
16	-0.104212	-0.137315	0.474613	0.056673	0.227176	1.2569e-08	1.246072
17	-0.104115	-0.137272	0.474607	0.056637	0.227143	6.9462e-09	1.481390
18	-0.104106	-0.137256	0.474611	0.056648	0.227141	5.5802e-09	1.169237
...							
28	-0.104268	-0.137209	0.474662	0.056843	0.227209	1.1516e-11	1.808904
29	-0.104264	-0.137210	0.474660	0.056839	0.227207	3.9831e-12	4.426169
30	-0.104262	-0.137210	0.474660	0.056836	0.227206	1.1488e-12	1.651176

In addition, if $f(\mathbf{s}^{k-1}) = f(\boldsymbol{\beta}^k)$, we have $\mathbf{s}^{k-1} = \boldsymbol{\beta}^k$ and $\boldsymbol{\beta}^k$ is an optimal solution; that is,

$$f(\boldsymbol{\beta}^k) = \min_{\boldsymbol{\beta}} f(\boldsymbol{\beta}). \tag{21}$$

Therefore, we have

$$\lim_{k\to\infty} f(\boldsymbol{\beta}^k) = \min_{\boldsymbol{\beta}} f(\boldsymbol{\beta}). \tag{22}$$

Proof $\boldsymbol{\beta}^k$ is computed by applying coordinate descent based on $\mathbf{s}^{k-1}$; that is,

$$\beta_i^k = \arg\min_{\beta} f([\beta_1^k, \ldots, \beta_{i-1}^k, \beta, s_{i+1}^{k-1}, \ldots, s_p^{k-1}]^T), \tag{23}$$

or equivalently

$$\beta_i^k = \frac{S(\mathbf{x}_i^T\mathbf{y} - \sum_{j<i}\mathbf{x}_i^T\mathbf{x}_j\beta_j^k - \sum_{j>i}\mathbf{x}_i^T\mathbf{x}_j s_j^{k-1}, \lambda)}{\|\mathbf{x}_i\|_2^2}. \tag{24}$$

Therefore, we have

$$\begin{aligned} &f([\beta_1^k, \ldots, \beta_{i-1}^k, \beta_i^k, s_{i+1}^{k-1}, \ldots, s_p^{k-1}]^T) \\ &\leq f([\beta_1^k, \ldots, \beta_{i-1}^k, s_i^{k-1}, s_{i+1}^{k-1}, \ldots, s_p^{k-1}]^T). \end{aligned} \tag{25}$$

for all i.

Since $\|\mathbf{x}_i\|_2 \neq 0$, $f([\beta_1^k, \ldots, \beta_{i-1}^k, \beta, s_{i+1}^{k-1}, \ldots, s_p^{k-1}]^T)$ is strongly convex in β. As a result, if the equality in Equation (25) holds, we have $\beta_i^k = s_i^{k-1}$. Recursively applying Equation (25), we have the following two facts: $f(\boldsymbol{\beta}^k) \leq f(\mathbf{s}^{k-1})$ and if $f(\boldsymbol{\beta}^k) = f(\mathbf{s}^{k-1})$, then $\mathbf{s}^{k-1} = \boldsymbol{\beta}^k$.

If $\mathbf{s}^{k-1} = \boldsymbol{\beta}^k$, it follows from Equation (24) that

$$\begin{aligned} \|\mathbf{x}_i\|_2^2\beta_i^k &= S(\mathbf{x}_i^T\mathbf{y} - \sum_{j\neq i}\mathbf{x}_i^T\mathbf{x}_j\beta_j^k, \lambda) \\ &= S(\mathbf{x}_i^T\mathbf{y} - \mathbf{x}_i^T X\boldsymbol{\beta}^k + \|\mathbf{x}_i\|_2^2\beta_i^k, \lambda), \end{aligned} \tag{26}$$

which leads to

$$\mathbf{x}_i^T\mathbf{y} - \mathbf{x}_i^T X\boldsymbol{\beta}^k \in \text{SGN}(\beta_i^k), \tag{27}$$

where

$$\text{SGN}(t) = \begin{cases} \{1\}, & t > 0 \\ \{-1\}, & t < 0 \\ [-1, 1], & t = 0. \end{cases} \tag{28}$$

Since $\boldsymbol{\beta}^*$ is an optimal solution to Equation (1) if and only if

$$\mathbf{x}_i^T\mathbf{y} - \mathbf{x}_i^T X\boldsymbol{\beta}^* \in \text{SGN}(\beta_i^*), \forall i, \tag{29}$$

it follows from Equation (27) that $\boldsymbol{\beta}^k$ is an optimal solution to Equation (1).

The relationship $f(\boldsymbol{\beta}^k) \geq f(\mathbf{s}^k)$ is guaranteed by the univariate optimization problem in Equation (12). Therefore, the sequence $\{f(\boldsymbol{\beta}^k)\}$ is decreasing. Meanwhile, the squence $\{f(\boldsymbol{\beta}^k)\}$ has a lower bound $\min_{\boldsymbol{\beta}} f(\boldsymbol{\beta})$. According to the well-known monotone convergence theorem, we have Equation (22).

This completes the proof of this theorem. □

4 Efficient Refinement Factor Computation

In this section, we discuss how to efficiently compute the refinement factor α^k in Equation (12). The function $g(\alpha)$ can be written as:

$$\begin{aligned} g(\alpha) &= \frac{1}{2}\left\|X((1-\alpha)\mathbf{h}^k + \alpha\boldsymbol{\beta}^k) - \mathbf{y}\right\|_2^2 + \lambda\|(1-\alpha)\mathbf{h}^k + \alpha\boldsymbol{\beta}^k\|_1 \\ &= \frac{1}{2}\left\|\mathbf{r}_h^k - \alpha(\mathbf{r}_h^k - \mathbf{r}^k)\right\|_2^2 + \lambda\|\mathbf{h}^k - \alpha(\mathbf{h}^k - \boldsymbol{\beta}^k)\|_1, \end{aligned} \tag{30}$$

where $\mathbf{r}_h^k = \mathbf{y} - X\mathbf{h}^k$ and $\mathbf{r}^k = \mathbf{y} - X\boldsymbol{\beta}^k$ are the residuals that correspond to $\mathbf{h}^k$ and $\boldsymbol{\beta}^k$, respectively. Note that 1) $\mathbf{r}_h^k = \mathbf{r}_s^{k-1}$ for SRRC and $\mathbf{r}_h^k = \mathbf{r}^{k-1}$ for SRRT, and 2) both $\mathbf{r}_h^k$ and $\mathbf{r}^k$ have been obtained before line 8 of Algorithm 2 and Algorithm 3. Before the convergence, we have $\mathbf{r}_h^k \neq \mathbf{r}^k$. Therefore, $g(\alpha)$ is strongly convex in α, and α^k, the minimizer to Equation (12), is unique.

When $\lambda = 0$, Equation (12) has a nice closed form solution,

$$\alpha^k = \frac{\langle \mathbf{r}_h^k, \mathbf{r}_h^k - \mathbf{r}^k \rangle}{\|\mathbf{r}_h^k - \mathbf{r}^k\|_2^2}. \tag{31}$$

Next, we discuss the case $\lambda > 0$. The subgradient of $g(\alpha)$ with regard to α can be computed as

$$\begin{aligned} \partial g(\alpha) = &\alpha\|\mathbf{r}_h^k - \mathbf{r}^k\|_2^2 - \langle \mathbf{r}_h^k, \mathbf{r}_h^k - \mathbf{r}^k \rangle \\ &+ \lambda \sum_{i=1}^{p} (\beta_i^k - h_i)\text{SGN}(h_i - \alpha(h_i - \beta_i^k)). \end{aligned} \tag{32}$$

Compute α^k is a root-finding problem. According to Theorem 1, we have $\alpha^k > 0$. Next, we consider only $\alpha > 0$ for $\partial g(\alpha)$. We consider the following three cases:

1. If $h_i = 0$, we have

$$(\beta_i^k - h_i)\text{SGN}(h_i - \alpha(h_i - \beta_i^k)) = \{|\beta_i^k|\}.$$

2. If $h_i(\beta_i^k - h_i) > 0$, we have

$$(\beta_i^k - h_i)\text{SGN}(h_i - \alpha(h_i - \beta_i^k)) = \{|\beta_i^k - h_i|\}.$$

3. If $h_i(\beta_i^k - h_i) < 0$, we let

$$w_i = \frac{h_i}{h_i - \beta_i^k}, \tag{33}$$

and we have

$$\begin{aligned} &(\beta_i^k - h_i)\text{SGN}(h_i - \alpha(h_i - \beta_i^k)) = \\ &\begin{cases} \{-|\beta_i^k - h_i|\} & \alpha \in (0, w_i) \\ \{|\beta_i^k - h_i|\} & \alpha \in (w_i, +\infty) \\ |\beta_i^k - h_i|\{[-1, 1]\} & \alpha = w_i. \end{cases} \end{aligned} \tag{34}$$

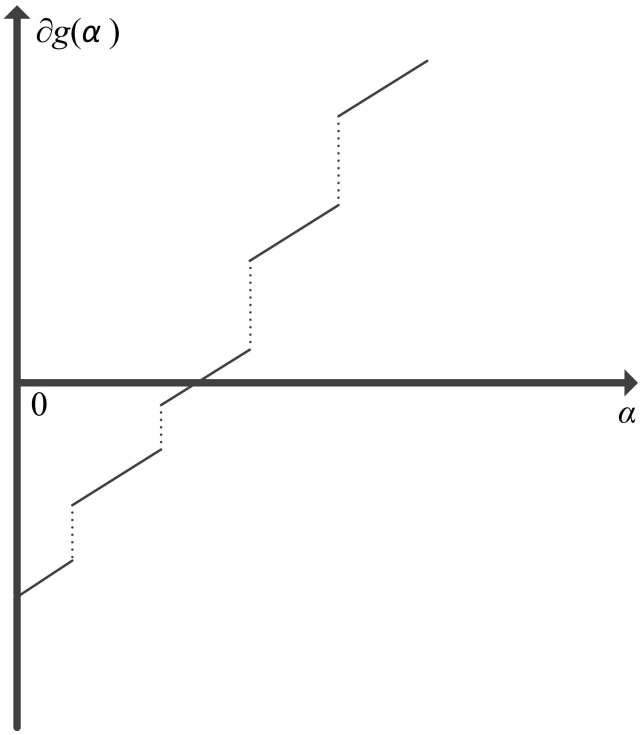

Figure 5: Illustration of $\partial g(\alpha)$. When $\lambda > 0$, it is a non-continuous piecewise linear monotonically increasing function. The intersection between $\partial g(\alpha)$ and the horizontal axis gives $\tilde{\alpha}^k$, the solution to Equation (12).

For the first two cases, the set SGN$(h_i - \alpha(h_i - \beta_i^k))$ is deterministic. For the third case, SGN$(h_i - \alpha(h_i - \beta_i^k))$ is deterministic when $\alpha \neq w_i$. Define

$$\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k) = \{i : h_i(\beta_i^k - h_i) < 0\}. \tag{35}$$

Figure 5 illustrates the function $\partial g(\alpha), \alpha > 0$. It can be observed that $\partial g(\alpha)$ is a piecewise linear function. If $\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)$ is empty, $\partial g(\alpha)$ is continuous; otherwise, $\partial g(\alpha)$ is not continuous at $\alpha = w_i, i \in \Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)$.

4.1 An Algorithm Based on Sorting

To compute the refinement factor, one approach is to sort w_i as follows:

First, we sort $w_i, i \in \Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)$, and assume $w_{i_0} \leq w_{i_1} \leq \ldots \leq w_{i_{|\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)|}}$.

Second, for $j = 1, 2, \ldots, |\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)|$, we evaluate $\partial g(\alpha)$ at $\alpha = w_{i_j}$ with the following three cases:

1. If $0 \in \partial g(w_{i_j})$, we have $\alpha^k = w_{i_j}$ and terminate the search.

2. If an element in $\partial g(w_{i_j})$ is positive, α^k lies in the piecewise line starting $\alpha = w_{i_{j-1}}$ and ending $\alpha = w_{i_j}$, and it can be analytically computed.

3. If all elements in $\partial g(w_{i_j})$ are negative, we set $j = j+1$ and continue the search.

Finally, if all elements in $\partial g(w_{i_j})$ are negative when $j = |\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)|$, α^k lies on the piecewise line that starts at $\alpha = w_{i_j}$. Thus, $\tilde{\alpha}^k$ can be analytically computed.

With a careful implementation, the naive approach can be completed in $O(p + m\log(m))$, where $m = |\Omega(\mathbf{h}^k, \boldsymbol{\beta}^k)|$. In Lasso, the solution is usually sparse, and thus m is much smaller than p, the number of variables.

4.2 An Algorithm Based on Bisection

A second approach is to make use of the improved bisection proposed in [7]. The idea is to 1) determine an initial guess of the interval $[\alpha_1, \alpha_2]$ to which the root belongs, where all elements in $\partial g(\alpha_1)$ are negative and all elements in $\partial g(\alpha_2)$ are positive, 2) evaluate $\partial g(\alpha)$ at $\alpha = \frac{\alpha_1+\alpha_2}{2}$ and update the interval to $[\alpha_1, \alpha)$ if all the elements in $\partial g(\alpha)$ are positive or to $[\alpha, \alpha_2)$ if all the elements in $\partial g(\alpha)$ are negative, 3) set the value of α to the largest value of w_i that satisfy $w_i < \alpha$ if all the elements in $\partial g(\alpha)$ are positive or to the smallest value of w_i that satisfy $w_i > \alpha$ if all the elements in $\partial g(\alpha)$ are negative, and 4) repeat 2) and 3) until finding the root of $\partial g(\alpha)$. With a similar implementation as in [7], the improved bisection approach has a time complexity of $O(p)$.

5 An Eigenvalue Analysis on the Proposed SRR

Let

$$A = XX^T = L + D + U, \tag{36}$$

where D is A's diagonal part, L is A's strictly lower triangular part, and U is A's strictly upper triangular part. It is easy to see that

$$L_{ij} = \begin{cases} \mathbf{x}_i^T\mathbf{x}_j & i < j \\ 0 & i \geq j, \end{cases} \tag{37}$$

$$D_{ij} = \begin{cases} \mathbf{x}_i^T\mathbf{x}_i & i = j \\ 0 & i \neq j, \end{cases} \tag{38}$$

$$U_{ij} = \begin{cases} \mathbf{x}_i^T\mathbf{x}_j & i > j \\ 0 & i \leq j. \end{cases} \tag{39}$$

We can rewrite Equation (2) as

$$D_{ii}\beta_i^k = S(\mathbf{x}_i^T\mathbf{y} - L_{i:}\boldsymbol{\beta}^k - U_{i:}\boldsymbol{\beta}^{k-1}, \lambda), \tag{40}$$

where $L_{i:}$ and $U_{i:}$ denote the ith row of L and U, respectively. Therefore, we can write coordinate descent iteration as:

$$D\boldsymbol{\beta}^k = S(X^T\mathbf{y} - L\boldsymbol{\beta}^k - U\boldsymbol{\beta}^{k-1}, \lambda). \tag{41}$$

When $\lambda = 0$, Equation (41) becomes

$$(L + D)\boldsymbol{\beta}^{k+1} = X^T\mathbf{y} - U\boldsymbol{\beta}^k, \tag{42}$$

which is the Gauss-Seidel method for solving

$$X^TX\boldsymbol{\beta} = (L + D + U)\boldsymbol{\beta} = X^T\mathbf{y}. \tag{43}$$

Equation (43) is also the optimality condition for Equation (1) when $\lambda = 0$. Our next discussion is for the case $\lambda = 0$ because it is easy to write the linear systems for the iterations.

Denote

$$G = -(L + D)^{-1}U. \tag{44}$$

Let G have the following eigendecomposition:

$$G = P\Delta P^{-1}, \tag{45}$$

where $\Delta = \text{diag}(\delta_1, \delta_2, \ldots, \delta_p)$ is a diagonal matrix consisting of its eigenvalues.

Lemma 1 *The magnitudes of the eigenvalues of G are all less than or equal to 1; that is,*

$$|\delta_i| \leq 1, \forall i. \tag{46}$$

Proof Let

$$G\mathbf{z} = \sigma\mathbf{z}, \tag{47}$$

where σ is an eigenvalue of G with the corresponding eigenvector being $\mathbf{z}$. Note that σ and the entries in $\mathbf{z}$ can be complex. Using Equation (36) and Equation (44), we have

$$(L + D - X^TX)\mathbf{z} = -U\mathbf{z} = (L + D)\sigma\mathbf{z}. \tag{48}$$

which leads to

$$(L + D)(1 - \sigma)\mathbf{z} = X^TX\mathbf{z}. \tag{49}$$

If $\sigma = 1$, the corresponding eigenvector $\mathbf{z}$ is in the null space of X^TX. If $\sigma \neq 1$, we have

$$(L + D)\mathbf{z} = \frac{1}{(1 - \sigma)}X^TX\mathbf{z}. \tag{50}$$

Premultiplying Equation (50) by $\mathbf{z}^H$, the conjugate transpose of $\mathbf{z}$, we have

$$\mathbf{z}^H(L + D)\mathbf{z} = \frac{1}{(1 - \sigma)}\mathbf{z}^HX^TX\mathbf{z}. \tag{51}$$

Taking the conjugate transpose of Equation (51), we have

$$\mathbf{z}^H(U + D)\mathbf{z} = \frac{1}{(1 - \bar{\sigma})}\mathbf{z}^HX^TX\mathbf{z}, \tag{52}$$

where $\bar{\sigma}$ denotes the conjugate of σ. Adding Equation (51) and Equation (52) and subtracting $\mathbf{z}^H X^T X \mathbf{z}$, we have

$$\mathbf{z}^H D \mathbf{z} = \left(\frac{1}{(1-\bar{\sigma})} + \frac{1}{(1-\sigma)} - 1 \right) \mathbf{z}^H X^T X \mathbf{z}. \tag{53}$$

Since $\mathbf{z}^H D \mathbf{z} > 0$ and $\mathbf{z}^H X^T X \mathbf{z} \geq 0$, we have

$$0 < \frac{1}{(1-\bar{\sigma})} + \frac{1}{(1-\sigma)} - 1 = \frac{1-|\sigma|^2}{|1-\sigma|^2}. \tag{54}$$

Therefore, we have $1 - |\sigma|^2 > 0$ or equivalently $|\sigma| < 1$.

This ends the proof of this lemma. □

5.1 An Eigenvalue Analysis on CD+SRRC

For CD+SRRC in Algorithm 2, when $\lambda = 0$ we have

$$\boldsymbol{\beta}^k = (L+D)^{-1}[X^T \mathbf{y} - U \mathbf{s}^{k-1}] \tag{55}$$

$$\mathbf{s}^k = (1-\alpha^k)\mathbf{s}^{k-1} + \alpha^k \boldsymbol{\beta}^k. \tag{56}$$

It can be shown that

$$\mathbf{s}^k - \mathbf{s}^{k-1} = \alpha^k \left[-(L+D)^{-1}U + \frac{1-\alpha^{k-1}}{\alpha^{k-1}} I \right] (\mathbf{s}^{k-1} - \mathbf{s}^{k-2}), \tag{57}$$

$$\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1} = -(L+D)^{-1}U(\mathbf{s}^{k-1} - \mathbf{s}^{k-2}). \tag{58}$$

When $k \geq 2$, we denote

$$A^k = \alpha^k \left[-(L+D)^{-1}U + \frac{1-\alpha^{k-1}}{\alpha^{k-1}} I \right]. \tag{59}$$

It can be shown that

$$A^k = P \Sigma^k P^{-1}, \tag{60}$$

where $\Sigma^k = \text{diag}(\sigma_1^k, \sigma_2^k, \ldots, \sigma_p^k)$ is a diagonal matrix and

$$\sigma_i^k = \alpha^k (\delta_i + \frac{1-\alpha^{k-1}}{\alpha^{k-1}}). \tag{61}$$

Therefore, we have

$$\mathbf{s}^k - \mathbf{s}^{k-1} = \left(\Pi_{i=2}^k A^k \right) (\mathbf{s}^1 - \mathbf{s}^0) = P(\Pi_{i=2}^k \Sigma_k) P^{-1} (\mathbf{s}^1 - \mathbf{s}^0), \tag{62}$$

For discussion convenience, we let $\Sigma^1 = \Delta$ and

$$T^k = \text{diag}(t_1^k, t_2^k, \ldots, t_p^k) = \Pi_{i=1}^k \Sigma^k, \tag{63}$$

where

$$t_i^k = \Pi_{i=1}^k \sigma_i^k. \tag{64}$$

We have

$$\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1} = PT^{k-1}P^{-1}(\mathbf{s}^1 - \mathbf{s}^0). \tag{65}$$

For the traditional coordinate descent in Algorithm 1, $\alpha^k = 1, \forall k$. For the proposed CD+SRRC in Algorithm 2, α^k optimizes Equation (12).

For the data used in Table 1, the eigenvalues of $-(L+D)^{-1}U$ are

$$\delta_1 = 0, \delta_2 = 0.00219338, \delta_3 = 0.12412229,$$

$$\delta_4 = 0.62606165, \delta_5 = 0.93956707.$$

For the traditional coordinate descent in Algorithm 1, when $k = 30$, we have

$$t_1^{29} = 0, t_2^{29} < 0.000001, t_3^{29} < 0.000001,$$

$$t_4^{29} = 0.000002, t_5^{29} = 0.164023.$$

For the coordinate descent with SRRC in Algorithm 2, we have

$$t_1^{29} = 0, t_2^{29} < 0.000001, t_3^{29} < 0.000001,$$

$$t_4^{29} < 0.000001, t_5^{29} = 0.000008.$$

This explains why the proposed SRRC can greatly accelerate the convergence of the coordinate descent method.

5.2 An Eigenvalue Analysis on CD+SRRT

For coordinate descent with SRRT in Algorithm 3, when $\lambda = 0$ we have

$$\boldsymbol{\beta}^k = (L+D)^{-1}[X^T\mathbf{y} - U\mathbf{s}^{k-1}] \tag{66}$$

$$\mathbf{s}^k = (1-\alpha^k)\boldsymbol{\beta}^{k-1} + \alpha^k\boldsymbol{\beta}^k. \tag{67}$$

When $k \geq 3$, it can be shown that

$$\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1} = G[(1-\alpha^{k-2})(\boldsymbol{\beta}^{k-2} - \boldsymbol{\beta}^{k-3}) + \alpha^{k-1}(\boldsymbol{\beta}^{k-1} - \boldsymbol{\beta}^{k-2})]. \tag{68}$$

When $k = 2$, we have

$$\boldsymbol{\beta}^2 - \boldsymbol{\beta}^1 = \alpha^1 G(\boldsymbol{\beta}^1 - \boldsymbol{\beta}^0). \tag{69}$$

Using the recursion in Equation (68), we can get

$$\boldsymbol{\beta}^3 - \boldsymbol{\beta}^2 = [(1-\alpha^1)G + \alpha^1 G\alpha^2 G](\boldsymbol{\beta}^1 - \boldsymbol{\beta}^0). \tag{70}$$

$$\begin{aligned}\boldsymbol{\beta}^4 - \boldsymbol{\beta}^3 =& [\alpha^1 G(1-\alpha^2)G + (1-\alpha^1)G\alpha^3 G \\ &+ \alpha^1 G\alpha^2 G\alpha^3 G](\boldsymbol{\beta}^1 - \boldsymbol{\beta}^0).\end{aligned} \tag{71}$$

Generally speaking, we can write

$$\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1} = PT^{k-1}P^{-1}(\boldsymbol{\beta}^1 - \boldsymbol{\beta}^0), \tag{72}$$

where $T^k = \text{diag}(t_1^k, t_2^k, \ldots, t_p^k)$ is a diagonal matrix. For t_i^k, it is a polynomial function of δ_i; that is, $t_i^k = \phi_k(\delta_i)$, where

$$\phi_k(t) = \underbrace{t \times \ldots \times t}_{\lceil k/2 \rceil} \sum_{i=0}^{k-\lceil k/2 \rceil} c_i \times \underbrace{t \times \ldots \times t}_{i}), \tag{73}$$

and $c_0, c_1, \ldots, c_{k-\lceil k/2 \rceil}$ are dependent on $\alpha^1, \alpha^2, \ldots, \alpha^{k-1}$. When $k = 2$, we have

$$c_0 = \alpha^1. \tag{74}$$

When $k = 3$, we have

$$\begin{array}{rcl} c_0 & = & 1 - \alpha^1, \\ c_1 & = & \alpha^1 \alpha^2. \end{array} \tag{75}$$

When $k = 4$, we have

$$\begin{array}{rcl} c_0 & = & \alpha^1(1 - \alpha^2) + \alpha^3(1 - \alpha^1), \\ c_1 & = & \alpha^1 \alpha^2 \alpha^3. \end{array} \tag{76}$$

When $k = 5$, we have

$$\begin{array}{rcl} c_0 & = & (1 - \alpha^2)(1 - \alpha^3), \\ c_1 & = & \alpha^1(1 - \alpha^2)\alpha^4 + \alpha^3(1 - \alpha^1)\alpha^4 + \alpha^1\alpha^2(1 - \alpha^3), \\ c_2 & = & \alpha^1 \alpha^2 \alpha^3 \alpha^4. \end{array} \tag{77}$$

For the coordinate descent with SRRT in Algorithm 3, we have

$$t_1^{29} = 0, t_2^{29} < 0.000001, t_3^{29} < 0.000001,$$

$$t_4^{29} = 0.000002, t_5^{29} = 0.000393,$$

which are smaller than the ones in the traditional coordinate descent shown in Section 5.1.

6 Related Work

In this section, we compare our proposed SRR with successive over-relaxation [17] and the accelerated gradient descent method [9].

6.1 Relationship between SRRC and Successive Over-Relaxation

Successive over-relaxation (SOR) is a classical approach for accelerating the Gauss-Seidel approach. Our discussion in this section considers only $\lambda = 0$, because SOR targets the acceleration of the Gauss-Seidel approach.

From Equation (43), we have

$$(wL + D)\boldsymbol{\beta} = wX^T\mathbf{y} - wU\boldsymbol{\beta} - (w - 1)D\boldsymbol{\beta}, \forall w > 0. \tag{78}$$

The iteration used in successive over-relaxation is:

$$\boldsymbol{\beta}^k = (wL + D)^{-1}[wX^T y - [wU + (w-1)D]\boldsymbol{\beta}^{k-1}, \tag{79}$$

which can be obtained by plugging $\boldsymbol{\beta}^k$ and $\boldsymbol{\beta}^{k-1}$ into Equation (78). Equation (79) can be rewritten as:

$$\boldsymbol{\beta}^k = \boldsymbol{\beta}^{k-1} - w(wL + D)^{-1}[X^T X\boldsymbol{\beta}^{k-1} - X^T\mathbf{y}]. \tag{80}$$

For the proposed CD+SRRC in Algorithm 2, when $\lambda = 0$ we have

$$\begin{aligned} \mathbf{s}^k &= (1-\alpha^k)\mathbf{s}^{k-1} + \alpha^k(L+D)^{-1}\left[X^T\mathbf{y} - U\mathbf{s}^{k-1}\right] \\ &= \mathbf{s}^{k-1} - \alpha^k(L+D)^{-1}\left[X^T X\mathbf{s}^{k-1} - X^T\mathbf{y}\right]. \end{aligned} \tag{81}$$

When $w = 1$ and $\alpha^k = 1$, both SOR and CD+SRRC reduce to the traditional coordinate descent. Equation (79) and Equation (81) share the following two similaries: 1) both make use of the gradient in the recursive iterations in that $X^T X\boldsymbol{\beta}^{k-1} - X^T\mathbf{y}$ is the gradient of $\frac{1}{2}\|X\boldsymbol{\beta} - \mathbf{y}\|_2^2$ at $\boldsymbol{\beta}^{k-1}$ and $X^T X\mathbf{s}^{k-1} - X^T\mathbf{y}$ is the gradient of $\frac{1}{2}\|X\mathbf{s}^{k-1} - \mathbf{y}\|_2^2$ at $\mathbf{s}^{k-1}$, and 2) both use a precondition matrix in that SOR uses $(wL + D)$ whereas SRRC uses $(L + D)$. A key difference is that the precondition matrix used in SRRC is parameter-free whereas the one used in SOR has a parameter. As a result, we can perform an inexpensive univariate search to find the optimal α^k used in SRRC whereas it is usually expensive for SOR to search for an optimal w in the same way as Equation (12).

When the design matrix has some special structures, it has been shown in [17] that the optimal value of w can be found for SOR. However, for the general design matrix X, it is hard to obtain the optimal w used for SOR. This might be a major reason that SOR is not widely used in solving Lasso with coordinate descent. For our proposed SRRC, the criterion in Equation (12) enables us to adaptively set the refinement factor α^k.

6.2 Relationship between SRRT and the Nesterov's Method

The SRRT scheme presented in Figure 4 is similar to the Nesterov's method in that both make use of a search point in the iterations. In addition, both set the search point using

$$\mathbf{s}^k = (1-\alpha^k)\boldsymbol{\beta}^{k-1} + \alpha^k\boldsymbol{\beta}^k. \tag{82}$$

However, the key difference is that the α^k used in the Nesterov's method is predefined according to a specified formula, whereas the α^k used in SRRT is set to optimize the objective function as shown in Equation (12). Note that if the Nesterov's method sets the α^k to optimize the objective function, it reduces to the traditional steepest descent method thus the good acceleration property of the Nesterov's method is gone.

Table 5: Performance for the synthetic data sets. The results are averaged over 10 runs. The sparsity is defined as the number of zeros in the solution divided by the total number of variables p.

data size	λ	CD	CD+SRRC	CD+SRRT	sparsity
$n = 500$	0.5	10.0	8.8	9.2	0.9395
	0.1	151.7	74.7	59.5	0.6406
$p = 1000$	0.05	463.2	179.0	109.1	0.5763
	0.01	4132.7	1419.4	326.1	0.5146
$n = 1000$	0.5	7.9	7.7	7.7	0.9397
	0.1	54.9	31.7	29.0	0.473
$p = 1000$	0.05	125.4	59.9	47.7	0.3126
	0.01	748.0	293.4	128.9	0.118
$n = 1000$	0.5	7.9	7.7	7.8	0.8856
	0.1	26.3	17.3	17.1	0.3102
$p = 500$	0.05	35.5	21.3	20.2	0.1678
	0.01	47.9	26.3	25.1	0.0382

Table 6: Performance for the real data sets. The sparsity is defined as the number of zeros in the solution divided by the total number of variables p.

data size	$\frac{\lambda}{\|X^T\mathbf{y}\|_\infty}$	CD	CD+SRRC	CD+SRRT	sparsity
$n = 38$	0.5	122	68	84	0.9982
	0.1	155	90	103	0.9964
$p = 7129$	0.05	254	119	127	0.9961
	0.01	2053	424	343	0.9948
$n = 62$	0.5	31	21	24	0.9975
	0.1	157	68	78	0.9840
$p = 2000$	0.05	308	115	118	0.9775
	0.01	2766	929	375	0.9715
$n = 6000$	0.5	26	16	11	0.9994
	0.1	180	103	108	0.9932
$p = 5000$	0.05	823	337	432	0.9876
	0.01	4621	1387	1368	0.8340

7 Experiments

In this section, we report experimental results for synthetic and real data sets, studying the number of iterations of CD, CD+SRRC and CD+SRRT for solving Lasso. The consumed computational time is proportional to the number of iterations.

Synthetic Data Sets We generate the synthetic data as follows. The entries in the $n \times p$ design matrix X and the $n \times 1$ response $\mathbf{y}$ are drawn from a Gaussian distribution. We try the following three settings of n and p: 1) $n = 500, p = 1000$, 2) $n = 1000, p = 1000$, and 3) $n = 1000, p = 500$.

Real Data Sets We make use of the following three real data sets provided in [2]: leukemia, colon, and gisette. The leukemia data set has $n = 38$ samples and $p = 7129$ variables. The colon data set has $n = 62$ samples and $p = 2000$ variables. The gisette data set has $n = 6000$ samples and $p = 5000$ variables.

Experimental Settings For the value of the regularization parameter, we try $\lambda = r\|X^T\mathbf{y}\|_\infty$, where $r = 0.5, 0.1, 0.05, 0.01$. For the synthetic data sets, the reported results are averaged over 10 runs. For a particular regularization parameter, we first run CD in Algorithm 1 until $\|\boldsymbol{\beta}^k - \boldsymbol{\beta}^{k-1}\|_2 \leq 10^{-6}$, and then run CD+SRRC and CD+SRRT until the obtained objective function value is less than or equal to the one obtained by CD.

Results Table 5 and Table 6 show the results for the synthetic and real data sets, respectively. The last column of each table shows the sparsity of the obtained Lasso solution, which is defined as the number of zero entries in the solution divided by the number of variables p. Figure 6 visualizes the results in these two tables. We can see that when the solution is very sparse (for example, $\lambda = 0.5\|X^T\mathbf{y}\|_\infty$), the proposed CD+SRRC and CD+SRRT consume comparable number of iterations to the traditional CD. The reason is that the optimal refinement factor computed by SRR in Equation (12) is equal to or close to 1, and thus CD+SRRC and CD+SRRT is very close to the traditional CD. Note that a regularization parameter $\lambda = 0.5\|X^T\mathbf{y}\|_\infty$ is usually too large for practical applications because it selects too few variables, and we usually need to try a smaller $\lambda = r\|X^T\mathbf{y}\|_\infty$ for example, $r = 0.01$. It can be observed that the proposed CD+SRRC and CD+SRRT requires much fewer iterations, especially for smaller regularization parameters.

8 Conclusion

In this paper, we propose a novel technique called successive ray refinement. Our proposed SRR is motivated by an interesting ray-continuation property on the coordinate descent iterations: for a particular coordinate, the value obtained in the next iteration almost always lies on a ray that starts at its previous iteration and passes through the current iteration. We propose two schemes for SRR and apply them to solving Lasso with coordinate descent. Empirical results for real and synthetic data sets show that the proposed SRR can significantly reduce the number of coordinate descent iterations, especially when the regularization parameter is small.

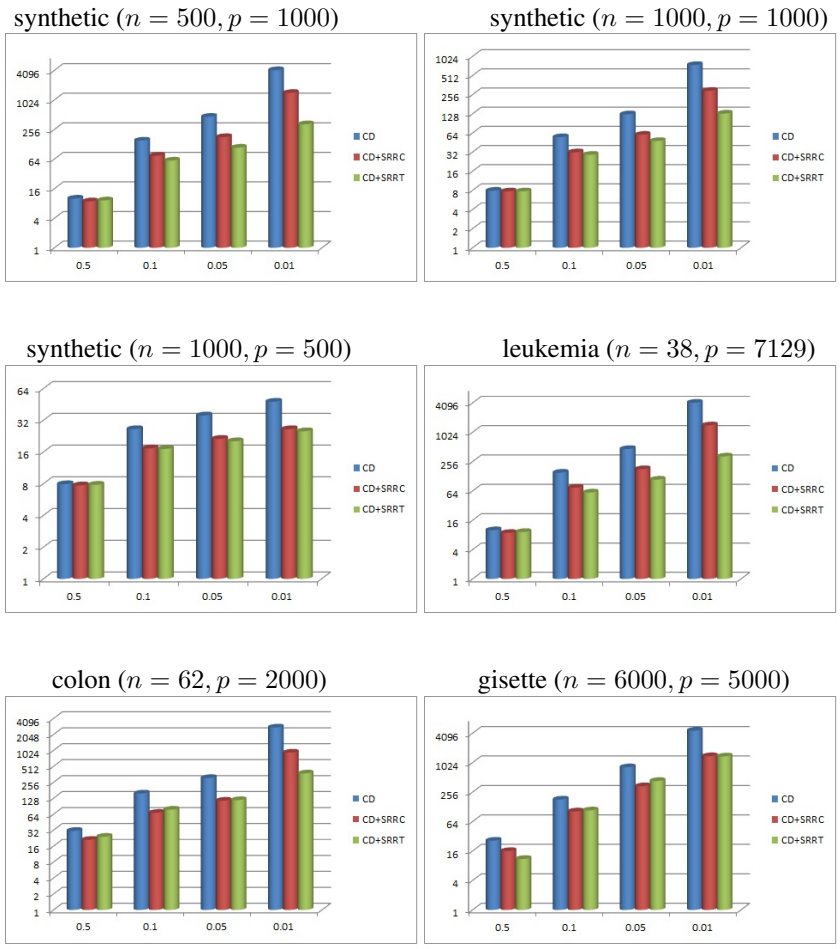

Figure 6: Number of iterations used by CD, CD+SRRC, and CD+SRRT on synthetic and real data sets. For all the plots, the x-axis corresponds to the regularization parameter $r = \frac{\lambda}{\|X^T \mathbf{y}\|_\infty}$ and the y-axis denotes the number of iterations in a logarithmic scale by different approaches.

We have established the convergence of CD+SRR, and it is interesting to study the convergence rate. We focus on a least squares loss function in (1), and we plan to apply the SRR technique to solving the generalized linear models. We compute the refinement factor as an optimal solution to Equation (12), and we plan to obtain the refinement factor as an approximate solution, especially in the case of generalized linear models.

References

[1] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

[2] C. C. Chang and C. J. Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011.

[3] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Annals of Statistics*, 32:407–499, 2004.

[4] J. H. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1):1–22, 2010.

[5] L. Ghaoui, V. Viallon, and T. Rabbani. Safe feature elimination in sparse supervised learning. *Pacific Journal of Optimization*, 8:667–698, 2012.

[6] K. Koh, S. Kim, and S. Boyd. An interior-point method for large-scale l1-regularized logistic regression. *Journal of Machine Learning Research*, 8:1519–1555, 2007.

[7] J. Liu and J. Ye. Efficient euclidean projections in linear time. In *International Conference on Machine Learning*, 2009.

[8] J. Liu, Z. Zhao, J. Wang, and J. Ye. Safe screening with variational inequalities and its application to lasso. In *International Conference on Machine Learning*, 2014.

[9] Y. Nesterov. *Introductory lectures on convex optimization : a basic course*. Applied optimization. Kluwer Academic Publ., 2004.

[10] K. Ogawa, Y. Suzuki, and I. Takeuchi. Safe screening of non-support vectors in pathwise SVM computation. In *International Conference on Machine Learning*, 2013.

[11] S. Shalev-Shwartz and A. Tewari. Stochastic methods for ℓ_1 regularized loss minimization. In *Proceedings of the 26th International Conference on Machine Learning*, 2009.

[12] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58:267–288, 1996.

[13] R. Tibshirani, J. Bien, J. H. Friedman, T. Hastie, N. Simon, J. Taylor, and R. J. Tibshirani. Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society: Series B*, 74:245–266, 2012.

[14] P. Tseng and Communicated O. L. Mangasarian. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109:475–494, 2001.

[15] J. Wang, B. Lin, P. Gong, P. Wonka, and J. Ye. Lasso screening rules via dual polytope projection. In *Advances in Neural Information Processing Systems*, 2013.

[16] S.J. Wright, R.D. Nowak, and M.A.T. Figueiredo. Sparse reconstruction by separable approximation. *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.

[17] D. M. Yong. *Iterative methods for solving partial difference equations of elliptical type*. PhD thesis, Harvard University, May 1950.

[18] G. X. Yuan, C. H. Ho, and C. J. Lin. An improved glmnet for l1-regularized logistic regression. *Journal of Machine Learning Research*, 13:1999–2030, 2012.

[19] J. X. Zhen, X. Hao, and J. R. Peter. Learning sparse representations of high dimensional data on large scale dictionaries. In *Advances in Neural Information Processing Systems*, 2011.